

TITLE OF THE INVENTION

Method for correlating behavior between two elements of a system to determine the presence of mutual interaction between the elements.

CROSS REFERENCE TO RELATED APPLICATIONS

This application claims the benefit of U.S. Provisional Patent Application Ser. No. 60/273,512, filed March 7, 2001, and entitled "Using traffic to discover network topology," incorporated herein by reference in its entirety.

FIELD OF THE INVENTION

The present invention relates to system behavior measurement and correlation in general, and applications therefor.

BACKGROUND OF THE INVENTION

The practical application of behavior correlation techniques is well known in many fields of science and engineering. For example, correlations between levels of automobile traffic and traffic accidents often serve as both a starting point and a benchmark in road safety programs, while correlations between advertising expenditure over the course of an advertising campaign and consumer purchases are invaluable in determining return on investment and providing other measures of effectiveness.

Whether behavior correlation techniques may be used to determine the presence of mutual interaction between two elements of a system should be examined. Some systems, such as computer networks, have elements that mutually interact, such as by communicating with each other. Computer networks, for example, typically have many network devices that are physically and logically interconnected via communication links. While information about the devices, such as their network address, is typically readily available, information regarding how the devices are physically interconnected must typically be derived manually or by employing automatic topology discovery techniques. Thus, were behavior correlation techniques successfully applied to determine the presence of mutual interaction between elements of a computer network, such techniques could be

used to determine how network devices are interconnected. Unfortunately, some systems, including computer networks, have mutually interacting elements whose behavior functions diverge over time with respect to one another, thus making behavior correlation difficult.

SUMMARY OF THE INVENTION

The present invention provides a method for correlating behavior between two elements of a system to determine the presence of mutual interaction between the elements.

In one aspect of the present invention there is provided a method for correlating behavior between two elements of a system to determine the presence of mutual interaction between the elements, the method including measuring the behavior of two elements of a system over time with respect to mutual interaction, thereby producing two behavior functions, expressing a plurality of constraints on a theoretical distance c between the behavior functions, and determining that the elements are behaving as mutually interacting elements where there exists an actual distance c satisfies the constraints.

In another aspect of the present invention the expressing step comprises maintaining an upper bound and a lower bound on the theoretical distance c , and the determining step comprises determining where the lower bound is smaller than or equal to the higher bound for any of the constraints.

In another aspect of the present invention the expressing step comprises expressing each of the constraints using at least two time-consecutive samples (a_n , a_{n+1}) of one of the functions and at least one sample (b_n) of the other of the functions that is time-intermediate the time-consecutive samples.

In another aspect of the present invention the expressing step comprises expressing each of the constraints as $(a_n - b_n) \leq c \leq (a_{n+1} - b_n)$.

In another aspect of the present invention the expressing step comprises expressing each of the distances using at least two time-consecutive samples (a_n , a_{n+1}) of one of the functions and at least one sample (b_n) of the other of the functions that is taken at the same time as one of the time-consecutive samples.

In another aspect of the present invention the expressing step comprises expressing each of the constraints as $(a_n - b_n) \leq c \leq (a_{n+1} - b_n)$.

In another aspect of the present invention there is provided a method for correlating behavior between two elements of a system to determine the presence of mutual interaction between the elements, the method including measuring the behavior of two elements of a system over time with respect to mutual interaction, thereby producing two behavior functions, expressing a plurality of constraints for a plurality of theoretical distances c_i between the behavior functions, and determining that the elements are behaving as mutually interacting elements where there exists a plurality of actual distances c_i that satisfies the constraints.

In another aspect of the present invention the expressing step comprises maintaining an upper bound and a lower bound on each of the plurality of theoretical distances c_i , and the determining step comprises determining where the lower bound is smaller than or equal to the higher bound for any of the constraints.

In another aspect of the present invention the expressing step comprises expressing each of the constraints using a plurality of samples a_n of one of the functions taken at times a'_1, a'_2, \dots, a'_n , a first plurality of time-consecutive samples b_n and a second plurality of time-consecutive samples b'_n of the other of the functions taken at times b'_1, b'_2, \dots, b'_n , where $a'_1 \leq b'_1 \leq a'_2 \leq b'_2 \leq \dots a'_n \leq b'_n$, and selecting each of the constraints from of a set of constraints defined by the pattern $a_1 - b_1 \leq c_1 \leq a_2 - b_1, a_2 - b_2 \leq c_2 \leq \min(a_2 - b_1 + P(b'_2 - b'_1), a_3 - b_2) a_3 - b_3 \leq c_3 \leq \min(a_2 - b_1 + P(b'_3 - b'_1), a_3 - b_2 + P(b'_3 - b'_2), a_4 - b_3)$.

In another aspect of the present invention there is provided apparatus for correlating behavior between two elements of a system to determine the presence of mutual interaction between the elements, the apparatus including means for measuring the behavior of two elements of a system over time with respect to mutual interaction, thereby producing two behavior functions, means for expressing a plurality of constraints on a theoretical distance c between the behavior functions, and means for determining that the elements are behaving as mutually interacting elements where there exists an actual distance c satisfies the constraints.

In another aspect of the present invention the expressing means comprises means for maintaining an upper bound and a lower bound on the theoretical distance c , and

the determining means comprises means for determining where the lower bound is smaller than or equal to the higher bound for any of the constraints.

In another aspect of the present invention the expressing means comprises means for expressing each of the constraints using at least two time-consecutive samples (a_n, a_{n+1}) of one of the functions and at least one sample (b_n) of the other of the functions that is time-intermediate the time-consecutive samples.

In another aspect of the present invention the expressing means comprises means for expressing each of the constraints as $(a_n - b_n) \leq c \leq (a_{n+1} - b_n)$.

In another aspect of the present invention the expressing means comprises means for expressing each of the distances using at least two time-consecutive samples (a_n, a_{n+1}) of one of the functions and at least one sample (b_n) of the other of the functions that is taken at the same time as one of the time-consecutive samples.

In another aspect of the present invention the expressing means comprises means for expressing each of the constraints as $(a_n - b_n) \leq c \leq (a_{n+1} - b_n)$.

In another aspect of the present invention there is provided apparatus for correlating behavior between two elements of a system to determine the presence of mutual interaction between the elements, the apparatus including means for measuring the behavior of two elements of a system over time with respect to mutual interaction, thereby producing two behavior functions, means for expressing a plurality of constraints for a plurality of theoretical distances c_i between the behavior functions, and means for determining that the elements are behaving as mutually interacting elements where there exists a plurality of actual distances c_i that satisfies the constraints.

In another aspect of the present invention the expressing means comprises means for maintaining an upper bound and a lower bound on each of the plurality of theoretical distances c_i , and where the determining means comprises means for determining where the lower bound is smaller than or equal to the higher bound for any of the constraints.

In another aspect of the present invention the expressing means comprises means for expressing each of the constraints using a plurality of samples a_n of one of the functions taken at times a'_1, a'_2, \dots, a'_n , a first plurality of time-consecutive samples b_n and a

second plurality of time-consecutive samples b'_n of the other of the functions taken at times $b'_1, b'_2 \dots, b'_n$, where $a'_1 \leq b'_1 \leq a'_2 \leq b'_2 \leq \dots a'_n \leq b'_n$, and means for selecting each of the constraints from of a set of constraints defined by the pattern $a_1 - b_1 \leq c_1 \leq a_2 - b_1$ $a_2 - b_2 \leq c_2 \leq \min(a_2 - b_1 + P(b'_2 - b'_1), a_3 - b_2)$ $a_3 - b_3 \leq c_3 \leq \min(a_2 - b_1 + P(b'_3 - b'_1), a_3 - b_2 + P(b'_3 - b'_2), a_4 - b_3)$.

The disclosures of all patents, patent applications, and other publications mentioned in this specification and of the patents, patent applications, and other publications cited therein are hereby incorporated by reference in their entirety.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated more fully from the following detailed description taken in conjunction with the appended drawings in which:

Fig. 1 is a simplified flowchart illustration of a method of behavior correlation, operative in accordance with a preferred embodiment of the present invention;

Fig. 2 is a simplified graphical illustration of behavior functions of two system elements, useful in understanding the method of Fig. 1;

Fig. 3 is a simplified flowchart illustration of a method of behavior correlation of mutually interacting system elements with diverging behavior functions, operative in accordance with a preferred embodiment of the present invention; and

Fig. 4 is a simplified graphical illustration of diverging behavior functions of two system elements, useful in understanding the method of Fig. 3.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Reference is now made to Fig. 1, which is a simplified flowchart illustration of a method of behavior correlation of mutually interacting system elements, operative in accordance with a preferred embodiment of the present invention, and additionally to Fig. 2, which is a simplified graphical illustration of behavior functions of two system elements, useful in understanding the method of Fig. 1. In the method of Fig. 1, the behavior of two elements of a system is measured over time with respect to mutual interaction. For example, in a computer network, the number of data packets transmitted by a device port A

is measured over time, as is the number of data packets received by a device port B over the same time period (although the counters used might start with different values). In the present embodiment the behavior of A and B is expressed as monotonic increasing functions. Samples $a_1, a_2 \dots, a_n$ are taken of port A at times $a'_1, a'_2 \dots, a'_n$, as are samples $b_1, b_2 \dots, b_n$ taken of port B at times $b'_1, b'_2 \dots, b'_n$, such that $a'_1 \leq b'_1 \leq a'_2 \leq b'_2 \leq \dots a'_n \leq b'_n$. A series of constraints on a theoretical distance c between the behavior functions of A and B may then be expressed using multiple $a_n-b_n-a_{n+1}$ sample groupings as follows:

$$(a_1 - b_1) \leq c \leq (a_2 - b_1)$$

$$(a_2 - b_2) \leq c \leq (a_3 - b_2)$$

...

$$(a_n - b_n) \leq c \leq (a_{n+1} - b_n)$$

Thus, for each constraint, at least two time-consecutive samples are taken from one of the behavior functions, and at least one sample is taken from the other of the behavior functions that is either time-intermediate the time-consecutive samples or that is taken at the same time as one of said time-consecutive samples. It will be appreciated that either of the behavior functions may serve as the source of a or b samples, provided that it is so used for each of the constraints.

If A and B are behaving as mutually interacting elements, then there exists an actual value c that satisfies all of the constraints. It may be seen that each constraint places an upper bound and a lower bound on c such that each inequality may act to tighten the bounds. Conversely, if there is no value c that satisfies the series of constraints, then it may be determined that A and B are not behaving as mutually interacting elements.

Some systems, including computer networks, have mutually interacting elements whose behavior functions diverge over time with respect to one another. For example, where a variable-length-packet switch mutually interacts with a fixed-length-cell switch, packets are converted into cells before transmission by the packet switch to the cell switch. Where a variable-length packet is not evenly divisible by the length of a fixed-length cell, the last cell is padded with zero bytes. Thus, a TCP/IP packet 1000 bytes in length is converted to 20 ATM cells of 56 bytes each, including an 8 byte header and 48 bytes of packet data, plus an additional ATM cell of 56 bytes in length, including an 8 byte

header, the last 40 bytes of packet data, and 8 bytes of padding. In order to measure the behavior of the two switches over time with respect to mutual interaction, the behavior functions of the switches should be expressed in common units of measure. However, since the packet switch counts only packets and bit octets (bytes) of actual data transmitted, while the cell switch counts only received cells which may include padding, multiplying the number of cells received at the cell switch by the number of bytes in a cell is likely to result in a received byte count that exceeds the transmitted byte count, thus resulting in diverging behavior functions, since subsequently transmitted packets might have padded cells as well.

Reference is now made to Fig. 3, which is a simplified flowchart illustration of a method of behavior correlation of mutually interacting system elements with diverging behavior functions, operative in accordance with a preferred embodiment of the present invention, and additionally to Fig. 4, which is a simplified graphical illustration of diverging behavior functions of two system elements, useful in understanding the method of Fig. 3. In the method of Fig. 3, the behavior of two elements of a system where diverging behavior may be present is measured over time with respect to mutual interaction. For example, in a computer network, where a determination is to be made of the mutual interaction between a cell switch A and a packet switch B, the number of data cells received by cell switch A is measured over time, as is the number of data packets and data octets transmitted by packet switch B over the same time period (although the counters used might start with different values). In the present embodiment the behavior of A and B are expressed as diverging monotonic increasing functions due to cell padding as described hereinabove. Sample cell counts a_1, a_2, \dots, a_n are taken of switch A at times a'_1, a'_2, \dots, a'_n , as are sample octet counts b_1, b_2, \dots, b_n and sample packet counts b'_1, b'_2, \dots, b'_n taken of switch B at times b'_1, b'_2, \dots, b'_n , such that $a'_1 \leq b'_1 \leq a'_2 \leq b'_2 \leq \dots \leq a'_n \leq b'_n$. It may then be determined that A and B are behaving as mutually interacting elements if there exists a series of distances c_1, c_2, \dots, c_n between the behavior functions of A and B such that:

$$1) a_1 \leq b_1 + c_1 \leq a_2 \leq b_2 + c_2 \leq \dots \leq b_n + c_n$$

and

$$2) \text{ for each } (i,j) \text{ where } i < j, c_j \leq c_i + P(b'_j - b'_i),$$

where P is a padding factor, such as 48 or 56 bytes in the present example,.

Thus, a series of constraints on theoretical distances $c_1, c_2, \dots c_n$ may be selected from of a set of constraints defined by the pattern below using multiple samples:

$$a_1 - b_1 \leq c_1 \leq a_2 - b_1$$

$$a_2 - b_2 \leq c_2 \leq \min(a_2 - b_1 + P(b'_2 - b'_1), a_3 - b_2)$$

$$a_3 - b_3 \leq c_3 \leq \min(a_2 - b_1 + P(b'_3 - b'_1), a_3 - b_2 + P(b'_3 - b'_2), a_4 - b_3)$$

...

For each constraint an upper bound and a lower bound is maintained on the distance c_i , where the samples are used to tighten the bound. Should the lower bound be larger than the higher bound for any of the constraints, the elements cannot mutually interact. Thus, if there exist actual distances c_i between the diverging monotonic increasing functions of A and B that satisfies the series of constraints, then it may be determined that A and B are behaving as mutually interacting elements. Otherwise, it may be determined that A and B are not behaving as mutually interacting elements.

It is appreciated that one or more of the steps of any of the methods described herein may be omitted or carried out in a different order than that shown, without departing from the true spirit and scope of the invention.

While the methods and apparatus disclosed herein may or may not have been described with reference to specific hardware or software, it is appreciated that the methods and apparatus described herein may be readily implemented in hardware or software using conventional techniques.

While the present invention has been described with reference to one or more specific embodiments, the description is intended to be illustrative of the invention as a whole and is not to be construed as limiting the invention to the embodiments shown. It is appreciated that various modifications may occur to those skilled in the art that, while not specifically shown herein, are nevertheless within the true spirit and scope of the invention.